Experimental Heat-Bath Cooling of Spins

G. Brassard[1], Y. Elias[2], J. M. Fernandez[3], H. Gilboa[4], J. A. Jones[5], T. Mor[2], Y. Weinstein[2] and L. Xiao[5]

[1] Département IRO, Université de Montréal, Montréal H3C 3J7, Canada
[2] Department of Computer Science, Technion, Haifa 32000, Israel
[3] Département de génie informatique, École Polytechnique de Montréal, Montréal H3C 3A7, Canada
[4] Department of Chemistry, Technion, Haifa 32000, Israel
[5] Centre for Quantum Computation, Clarendon Laboratory, University of Oxford, Parks Road, Oxford OX13PU, United Kingdom

arXiv:1404.6885v1 [quant-ph] 28 Apr 2014

Abstract

Algorithmic cooling (AC) is a method to purify quantum systems, such as ensembles of nuclear spins, or cold atoms in an optical lattice. When applied to spins, AC produces ensembles of highly polarized spins, which enhance the signal strength in nuclear magnetic resonance (NMR). According to this cooling approach, spin-half nuclei in a constant magnetic field are considered as bits, or more precisely, quantum bits, in a known probability distribution. Algorithmic steps on these bits are then translated into specially designed NMR pulse sequences using common NMR quantum computation tools. The *algorithmic* cooling of spins is achieved by alternately combining reversible, entropy-preserving manipulations (borrowed from data compression algorithms) with *selective reset*, the transfer of entropy from selected spins to the environment. In theory, applying algorithmic cooling to sufficiently large spin systems may produce polarizations far beyond the limits due to conservation of Shannon entropy.
Here, only selective reset steps are performed, hence we prefer to call this process "heat-bath" cooling, rather than algorithmic cooling. We experimentally implement here two consecutive steps of selective reset that transfer entropy from two selected spins to the environment. We performed such cooling experiments with commercially-available labeled molecules, on standard liquid-state NMR spectrometers. Our experiments yielded polarizations that *bypass Shannon's entropy-conservation bound*, so that the entire spin-system was cooled. This paper was initially submitted in 2005, first to Science and then to PNAS, and includes additional results from subsequent years (e.g. for resubmission in 2007). The PostScriptum includes more details.

1 Introduction

The limiting factor in Nuclear Magnetic Resonance (NMR) spectroscopy and imaging (MRI) is often the intrinsically low polarization of nuclear spins, which leads to low signal-to-noise ratio (SNR). Enhancement of the SNR would permit more rapid data acquisition, enabling more efficient analysis of chemicals and visualization of tissues. Alternatively, less material would be required, leading to reduced toxicity. Possible methods to increase the sensitivity of NMR [1, 2] include higher magnetic fields which are currently limited to about 20 tesla, signal averaging which is time consuming and temperature reduction, which are impractical for various biomedical applications. Another solution is *effective-cooling*; cooling the spins but not their environment, so that their *effective temperature* is lower than the temperature of the surrounding heat-bath. A spin-half particle in a constant and homogeneous magnetic field has a steady-state polarization bias, ε, inversely related to the temperature, T. Therefore, spins with polarization biases above their equilibrium bias are considered *cool*, relative to the temperature of the environment [2]. Polarization enhancement by a factor of M improves the SNR by the same factor. For the case relevant here, $\varepsilon \ll 1$, and then it is inversely *proportional* to the temperature.

Effective-cooling of spins, also called "spin cooling", is a well established concept. Many NMR studies rely on a variety of methods to transfer polarization among spins to increase signal intensity; see [2, 3] and references therein. The simplest polarization transfer (PT) method, commonly used in NMR, is a transfer of polarization from a more polarized (sensitive) spin to a significantly less polarized (insensitive) spin, e.g., from hydrogen (^{1}H) to carbon (^{13}C). Other (less common) PT methods, such as dynamic nuclear polarization, transfer polarization from electron spins to nuclear spins. Several enhanced effective-cooling methods have succeeded in creating very high polarizations by what one may call the "deep-freeze" approach, in which extreme physical cooling of the relevant spins is involved. These include dynamic nuclear polarization enhanced with physical cooling and heating [4], *para*-hydrogen in two-spin systems [5, 6] and hyperpolarized xenon [7–9]. We briefly describe various effective-cooling methods in Appendix A.

Two other approaches, suggested by Sørensen (1989) and by Schulman and Vazirani (1999), are fundamentally different from the above. They are based on general unitary transformations [3] and on closely related data compression methods in closed systems [10]. Both approaches use reversible entropy-preserving manipulations to cool specific spins, while heating others. These approaches are distinct from the simple PT and the deep-freeze approaches, as they allow cooling beyond the capacity of the high-bias spins. However, those entropy manipulation approaches in closed systems [3, 10] are limited by Shannon's entropy-conservation bound [11, 12] and Sørensen's unitarity bound [3]. The connection between data compression, entropy manipulations, and the resulting increased polarization bias is explained in [10, 13, 14]; see also Appendix B.1.

A more general approach, *algorithmic cooling* (AC), introduced by Boykin, Mor, Roychowdhury, Vatan and Vrijen, in 2002, makes use of entropy manipulations in open systems [13]. AC cools beyond the reversible schemes [3, 10] described above. In particular, when applied to spins, AC can cool beyond Shannon's entropy bound on closed-system entropy manipulations (the source coding theorem [11, 12]).

The entropy manipulation approaches in closed systems [10] and open systems [13] were originally suggested as methods for conceptually resolving (from a computational complexity point of view) the scalability problem of NMR Quantum Computing (NMRQC). The open-system method, AC, was later recognized also for its potential usefulness in NMR spectroscopy [14, 15], since moderate effective-cooling beyond Shannon's entropy bound is theoretically achievable even when AC is applied onto small molecules.

2 Cooling of Spins Beyond Shannon's Entropy Bound

From an information theoretic point of view, that is, when arbitrary closed-system entropy-preserving manipulations are allowed [3, 10], the limitations due to Shannon's bound on closed-system entropy manipulations can be shown succinctly via a simple example[1]: Consider our experimental 3-spin system, trichloroethylene, and recall that we assume polarization biases much smaller than one. This molecule contains three relevant spins; two labeled carbons with nearly identical equilibrium polarizations, and one hydrogen spin with an equilibrium polarization that is nearly four-fold higher. For simplicity, we take these polarizations in this section as ε for both carbons and 4ε for the hydrogen spin. The Shannon entropy of the system at thermal equilibrium is calculated (to leading order in ε) by summing the entropy of the three spins, $H_\varepsilon + H_\varepsilon + H_{4\varepsilon}$, where $H_\varepsilon \triangleq -\frac{1+\varepsilon}{2}\log_2\left(\frac{1+\varepsilon}{2}\right) - \frac{1-\varepsilon}{2}\log_2\left(\frac{1-\varepsilon}{2}\right) \xrightarrow{\varepsilon \ll 1} 1 - \varepsilon^2/\ln 4$, such that $H = 3 - (1^2 + 1^2 + 4^2)\varepsilon^2/\ln 4 \longrightarrow 3 - 18\varepsilon^2/\ln 4$, in *bit* units. The *information content* of the molecule, that is the difference from maximal entropy, is given in this case by $3 - H$, such that to very good approximation, in units of $\varepsilon^2/\ln 4$, the initial information content is

$$I_{initial}^{approx} \approx 18 \ . \tag{1}$$

In this paper we are mainly interested in increasing the value of I, thus decreasing the total entropy (H) and cooling the entire system. Bypassing Shannon's entropy bound on this 3-spin system means increasing the information content of the system above 18. This could be done, for instance, by increasing the polarization of a single spin above the value of $\sqrt{18}\varepsilon$ (see the Post-Scriptum for recent experimental results achieving that goal). It could also be done by reaching biases above the value of $\sqrt{6}\varepsilon$ on the three spins, or reaching biases of 2 on two spins and a bias above $\sqrt{10}$ on the third.

AC was introduced [13] in order to increase spin polarization beyond Shannon's entropy bound, and, for long molecules, even *far beyond* it (at least in theory). Notably, thermalization is used beneficially as an integral part of the cooling scheme, while ordinarily it is perceived as a major obstacle in quantum computation and NMRQC. AC employs slow-relaxing spins named *computation spins* and rapidly relaxing spins named *reset spins*, to cool the entire spin-system by pumping entropy to the environment. The ratio $\mathcal{R}$, between the spin-lattice relaxation (thermalization) times of the computation spins and the reset spins, must satisfy $\mathcal{R} \gg 1$, to permit the application of many cooling steps to the system, while the computation spins are still quite isolated from the environment, namely from the heat-bath. In principle, here are the three basic operations of AC [13, 14]: a.— COMPRESSION. Reversible entropy manipulation steps (on more than two spins) that redistribute the entropy in the system so that some spins become cooler than the environment, while other spins become hotter [3, 10]. Saying that certain spins became cooler implies that they became more polarized. b.— SWAP. Controlled interactions allow specific computation spins to adiabatically lose their entropy to reset spins that have lower entropy. This means that the polarization is transferred from the reset spins onto these computation spins. c.— WAIT. The reset spins rapidly thermalize, transferring their entropy to the environment, while the computation spins remain colder, so that the entire system is cooled.

This combined set of operations increases spin polarization and bypasses Shannon's entropy bound. AC uses these operations recursively, in order to obtain, theoretically, extremely low spin temperatures. Note that although the algorithms are classical (namely, can be defined using only classical bits), the implementations via spins make use of the tools developed in NMRQC such as the use of specific quantum gates (see reviews in [16, 17]). In particular, a universal set of gates (by definition) can be used to compose any algorithm, yet a subset might also be sufficient for running a specific algorithm such as AC.

With sufficiently large relaxation-times ratio ($\mathcal{R} \gg 1$) and identical initial bias (ε) for all computation and reset spins, AC can theoretically cool exponentially more than the closed-system entropy-preserving method: In 2004, Fernandez, Lloyd, Mor and Roychowdhury [14] designed cooling algorithms, named practicable AC (PAC) that use simple quantum gates, and that can be applied to *short molecules*; As long as the final bias is small[2], $\varepsilon_{final} \ll 1$, the application of a simple algorithm to $n-1$ computation spins and one reset spin (n is odd here, in this example) theoretically improves the polarization of a single computation spin by an exponential factor, yielding $\varepsilon_{final} \approx (3/2)^{(n-1)/2}\varepsilon$, while Shannon's bound restricts reversible cooling to $\varepsilon_{final} \approx \sqrt{n}\varepsilon$; see Appendix B, sections B.1 and B.2.

AC can (in theory) yield interesting nontrivial cooling, even on our simple 3-spin system described earlier, provided that the (hydrogen) spin with the higher initial polarization of 4ε is *also* a reset spin. In that case, AC can ideally increase the polarization of one carbon spin by a factor of 6 within very few steps, and asymptotically (in the limit of an infinite number of steps [19]) by a factor of 8. In comparison, PT can at most increase the polarization of one carbon spin by a factor of 4, and the optimal closed-system cooling is limited to cooling one spin by a factor of $\sqrt{18}$. The information content of our 3-spin system can (ideally) be increased using very simple AC protocols to 48 or even 68 within very few steps, and asymptotically to 96. For the derivations of those numbers, see Appendix B.3.

In reality, finite thermalization time ratios make cooling beyond Shannon's bound challenging, especially if one prefers to use conventional liquid-state NMR systems and commercially-available molecules. Realistic AC suffers from small $\mathcal{R}$ between the relaxation times (T_1) of relevant computation and reset spins ($T_1{}^{comp}$ and $T_1{}^{reset}$, respectively), and the conditions under which exponential cooling is still possible are not yet fully understood. Taking into consideration this gap and others (that are described later) between theory and experiment, it is not at all clear whether AC can yield the desired advantages even when applied to short molecules. In particular, for liquid-state NMR, the typical range of R in commercially-available candidate molecules is around 2-20.

3 Heat-bath Cooling of Spins

We apply here a simplified variant of AC, which we call "heat-bath cooling", to the 3-spin system trichloroethylene, shown in Figure 1; we experimentally bypass Shannon's entropy bound with the spin-system of this organic molecule.

[1] For a different example [14] see Appendix B.1.

[2] AC is good also for final biases approaching 1 [18].

Our results, initially appearing in the public domain in 2005 [20], have encouraged later research in this direction, which could eventually lead to applications already in the near future, as even moderate cooling can be found useful in various bio-medical NMR applications. This is in contrast with other applications of quantum computing, such as Shor's factorization algorithm, that probably cannot become useful in the near future.

To define heat-bath cooling, let us look again at the steps that compose algorithmic cooling. The three steps of AC can be combined together (interlaced) in several ways, yielding slightly modified definitions of the building blocks composing AC. For instance, compression steps can be applied directly on computational spins and reset spins together to compress entropy onto one specific reset spin [14, 18, 19]. Seen this way, AC is composed of (general) reversible entropy-preserving manipulations (containing COMPRESSION and SWAP steps as special cases) combined with WAIT steps.

Alternatively, one may view AC as being composed of COMPRESSION (or more generally, reversible entropy-preserving manipulations) and *selective reset* steps, in which the entropy is transferred from *selected* computation spins to the heat-bath. Each such selective-reset step is composed solely of SWAP and WAIT steps, aimed at cooling a specific computation spin by transferring its entropy to a specific (and more polarized) reset spin, and then waiting. Consequently, that reset spin thermalizes and conveys its excess entropy irreversibly to the heat-bath. The refreshed reset spin can then be reused for additional selective resets, e.g. to cool spins that are heated in a subsequent polarization compression. Each selective reset step thus potentially cools the entire spin system.

When selective-reset steps are performed, with no compression step(s), in order to cool computation spins, we refer to this degenerate version of AC as "heat-bath cooling". Here we experimentally show controlled entropy extraction from our 3-spin system via a *dual-selective-reset* process – two selective-reset steps, using a single reset bit; we thus perform the irreversible part of AC. The dual-selective-reset process is composed of four steps: 1.– transfer polarization from H to the far carbon (C1, see Figure 1); 2.– wait for a suitable amount of time, t_1, for H to repolarize; 3.– transfer polarization from H to the adjacent carbon (C2); 4.– wait an additional period of time, t_2, for H to repolarize. Our additional experimental goal here is to bypass Shannon's bound regarding conservation of the total entropy of our 3-spin system, via heat-bath cooling.

In the ideal case, starting from equilibrium biases, approximately $\{\varepsilon, \varepsilon, 4\varepsilon\}$, denoted $\{1, 1, 4\}$ for C1, C2 and H, respectively, our algorithm produces the following sequence of polarizations:

$$\begin{aligned} \{1,1,4\} &\xrightarrow{\text{SWAP(C1,H)}} \{4,1,1\} \xrightarrow{\text{WAIT}} \{4,1,4\} \\ &\xrightarrow{\text{SWAP(C2,H)}} \{4,4,1\} \xrightarrow{\text{WAIT}} \{4,4,4\} \end{aligned} \tag{2}$$

If, for instance, the initial temperature is $300K$, these final polarization biases correspond to final effective temperatures of 75K for each carbon and equilibrium (300K) for the proton. The final information content resulting from these final polarizations is

$$I_{final}^{approx} \approx 4^2 + 4^2 + 4^2 = 48 \ , \tag{3}$$

a nearly three-fold increase, which clearly bypasses the entropy-conservation bound.

4 Experimental Heat-bath Cooling

Our cooling experiments were performed on standard liquid-state NMR spectrometers with commercially available $^{13}C_2$-trichloroethylene (labeled TCE), shown in Figure 1, in which the hydrogen functions as a reset bit (reset spin) owing to its relatively rapid relaxation, and the two (labeled) carbons serve as the computation bits (computation spins). See Subsection 4.1 for more details about materials and methods used in our experiments.

From an algorithmic point of view, a transfer of polarization can be achieved by exchanging the states of the two spins, using a SWAP gate as in Eq. (2). The *required* transfer of polarization, however, is uni-directional as we are not concerned with residual polarization (after PT) on the hydrogen, because it is to be reset, regaining most of its initial polarization. Therefore, the implementation of PT is potentially simpler than that of a full (bi-directional) SWAP, as we explicitly show in Appendix C. Also note that it is often inefficient to directly implement PT between non-adjacent spins due to weak scalar couplings. Step 1 above thus comprises two sequential steps: 1a.– PT(H $\rightarrow$ C2); 1b.– PT(C2$\rightarrow$ C1). We refer to any pulse sequence implementing nontrivial heat-bath cooling (leading to cooling of at least two computation spins using one reset spin) via PT, and WAIT (in our case, steps 1a, 1b, 2, 3, and 4) as POTENT: POlarization Transfer via ENvironment Thermalization. In our POTENT experiment, we implement uni-directional PT using a variant of the INEPT (Insensitive Nuclear Enhancement by Polarization Transfer) pulse sequence [21]. Our pulse sequence uses exclusively (nonselective) "hard pulses", namely very short pulses over a wide frequency range, such that all nuclei of the same species (in our case, the two carbons) are similarly affected.

To estimate the required WAIT times, t_1 and t_2, one must account for the finite ratios of T_1 values. We performed a numerical simulation of the POTENT pulse sequence using a standard relaxation model and T_1 relaxation values which we measured in the laboratory. The simulation provides an estimated range where one should experimentally test for the optimal values of these two H repolarization delays, namely, delays that yield maximal final information content. For details of the simulation see section 6.

4.1 Materials and methods

Our POTENT experiments were performed on standard 400-600 MHz liquid-state NMR spectrometers in three different labs, using different solvents. The first successful experiment was performed at the Université de Montréal in March 2002 on a Bruker DMX-400 by the authors GB, JMF, TM and YW, together with Raymond Laflamme from University of Waterloo (Ontario, Canada). Initial and final carbon spectra of the original experiment are shown in Figure 3. Only ^{13}C spectra were recorded, therefore proton polarizations were calculated according to the simulation model. The calculated information content was beyond the entropy bound.

Many POTENT experiments were carried out later on at the Technion and at Oxford university, with various delays, t_1 and t_2. The results presented in the "Results" section were obtained at the Technion using a Bruker Avance 500 MHz spectrometer with deuterated chloroform ($CDCl_3$) as the solvent. $^{13}C_2$-trichloroethylene (TCE) was obtained from CDN Isotopes (99.2% ^{13}C) or from Cambridge Isotope Laboratories (99% ^{13}C, diisopropylamine stabilized). In Montréal and in Haifa, the same Bruker pulse sequences were used on samples of TCE in deuterated chloroform (Aldrich, 99.9%D). At Oxford University, a Varian Inova 600 MHz spectrometer was used; functionally equivalent pulse programs were applied to TCE samples in either deuterated chloroform or deuterated acetone.

Most experiments were carried out with a paramagnetic relaxation reagent, Cr(III)acetlyacetonate ($Cr(acac)_3$), obtained from Alfa Aesar (97.5+% pure), at a final concentration of about 0.2 mg/mL. The reagent was added to increase the relaxation times ratios, as suggested in [22]; Although Shannon's entropy bound was bypassed in both cases (with and without reagent), the experimental success was much more robust when the reagent was used. The results presented in the "Results" section were obtained with the relaxation reagent.

For implementing the PT, high-power, non-selective ("hard") pulses were applied, where the proton and the adjacent carbon, C2 (see Fig. 1) were on resonance. Selective addressing of one of the two ^{13}C spins was achieved by inducing a phase separation using the chemical shift, while refocusing the scalar coupling evolutions. More details can be seen in Appendix C.1. Signal averaging (and phase cycling) were not employed, in order to isolate the effect of heat-bath cooling.

5 Results

We acquired spectra of TCE at equilibrium and after the cooling pulse sequence. Figure 2 displays ^{13}C and ^{1}H NMR spectra for TCE. Figures 2(a) and 2(c) were obtained at thermal equilibrium and serve as a reference point for ^{13}C and ^{1}H, respectively. Figures 2(b) and 2(d) were obtained after the cooling pulse sequence. For both ^{13}C nuclei, the increase in polarization bias can be intuitively observed by looking at the noticeably higher peaks compared with the reference spectrum, while the ^{1}H intensity is only slightly reduced. As the polarization bias is directly proportional to the area under the relevant peak, its increase/decrease is more accurately calculated via integration. The experimentally acquired resonance frequencies (see the caption of Figure 1) at room temperature (296 ± 1 K) were used to calculate equilibrium biases of 1.000 ± 0.003 for the carbons and 3.98 ± 0.01 for the proton. The errors are mainly due to the temperature uncertainty, see Appendix C. These room-temperature polarization biases lead to an initial information content of

$$I_{initial}^{actual} = 17.8 \pm 0.1 \ . \tag{4}$$

A maximal bias identical to the hydrogen initial bias (3.98) could be achieved for all three spins if the T_1 ratios were infinite and all operations were ideal, leading to an ideal final information content of

$$I_{final}^{ideal} = 47.5 \pm 0.1 \ . \tag{5}$$

These values replace the approximate values presented in Eqs. (1) and (3), respectively. Both are calculated using the sum-of-squares of the relevant biases.

The spectra after the cooling and the final experimental polarization biases (presented in Table 1) were obtained with the delays $t_1 = 8$s and $t_2 = 12$s. In the numerical simulation that we performed, these delays reside on a plateau, in which a wide range of delay combinations give (more or less) the same highest information content. The specific values of $t_1 = 8$ and $t_2 = 12$ were obtained by experimentally optimizing the information content while also minimizing the overall duration of the POTENT pulse sequence.

For each nucleus, the *experimental* final polarization bias was obtained by comparing its integrals before and after the POTENT pulse sequence. The resulting final biases are $\{1.74 \pm 0.01, 1.86 \pm 0.01, 3.77 \pm 0.01\}$, corresponding to the final spin-temperatures given in Table 1; The small experimental error in the biases was obtained by repeating the experiment five times under the same conditions. Calculating the final bias of a specific spin from its integrals is equivalent to tracing out [23] the other two spins. The resulting calculation of the information content, in general, does not provide the information content, I, but a lower bound on it[3], which we denote as $\tilde{I}$. Such a calculation provides directly the information content (I) only when the state of the system constitutes a *tensor product state* of the three spins, as in the case above of the initial information content, and as in the cases of the two simulations described below.

The lower bound on the final experimental information content is

$$\tilde{I}_{final}^{actual} = 20.7 \pm 0.1 \ . \tag{6}$$

This presents an increase of about $16\% \pm 1\%$ over the information content at thermal equilibrium, 17.8, proving that Shannon's entropy bound was indeed bypassed.

The ideal final value of 47.5 is devoid of relaxation constraints, and therefore does not provide a reasonable prediction of the experimental information content. To obtain a better prediction of the experimental value we performed numerical simulations that take into account the experimental T_1 values (presented in the caption of Figure 1). Assuming perfect PT and our experimental delays ($t_1 = 8$s and $t_2 = 12$s), the simulated information content is about $I_{final}^{sim} \approx 29$. The remaining discrepancy between actual and simulated values is still large; this discrepancy can largely be ascribed to the low efficiency of the PT steps. The experimental PT efficiencies were about 92% in step 1a, 69% in step 1b, and 74% in step 1c (see Appendix C for details). A *practical* simulation[4], which goes beyond the ideal simulation by taking these imperfect PT efficiencies into account yields

$$I_{final}^{prac-sim} \approx 22 \ , \tag{7}$$

[3] It is a lower bound due to the subadditivity of Shannon entropy.

[4] Both the ideal simulation and the more practical simulation could be applied to other molecules.

which is much closer to the obtained experimental result. We estimate that the low PT efficiencies are mainly due to off-resonance effects, decoherence (dephasing), and imperfections in the pulse sequence. The ideal and practical simulations, as well as a few subtle factors that potentially contribute to the small remaining discrepancy between $I_{final}^{prac-sim}$ and $\tilde{I}_{final}^{actual}$, are described in section 6.

As we see in Table 1, both carbons of TCE were cooled considerably, well below 200K, following the application of POTENT. Note that, in addition to the experimental final biases, the table also includes the ideal and the practical simulated (final) biases.

If the goal of the effective-cooling is to reach minimal carbon spin-temperatures then the final thermalization period, namely step 4 of our heat-bath cooling, should be omitted (t_2 set close to 0). The resulting final temperatures of the far carbon (C1) and the adjacent carbon (C2) were 145 ± 2K (C1), and 101 ± 1K (C2). For more details on this experiment see Appendix C.2.

6 Simulations

We simulated heat-bath cooling experiments with Matlab (The MathWorks, Natick, MA, USA). During the two delays, t_1 and t_2, each spin is assumed to relax according to its experimental T_1. If SWAP gates were used for the polarization transfer, then the quantum states of the two spins involved in a SWAP are simply swapped. However, when dual-CNOT is used (as explained in Appendix C.1) polarization of the source spin is transfered to the target spin, and the polarization of the target spin is transfered into classical correlations, so that the resulting polarization of the source spin is zero after the transfer.

In the simulation we assume that indeed the resulting polarization of the source spin is zero after the transfer. However, correlation terms between the proton and the carbons were ignored on the basis of an extended-Markov[5] assumption: the proton was considered to some extent as part of the memory-less environment, because t_1 and t_2 are sufficiently large compared to T_1(H). On the one hand, its correlations with the carbons are assumed to fully cancel out during the reset, but on the other hand, its final polarization bias after the reset is calculated precisely taking into account the exact reset time. Correlation terms among the two carbons (after the first PT from H to C1) cancel out due to the reset of the near-carbon with the help of the uncorrelated proton.

The polarization biases were calculated as a function of time by the formula

$$\varepsilon^S(t^S) = (\varepsilon_{init}^S - \varepsilon_{eq}^S)e^{-\frac{t^S}{{T_1}^S}} + \varepsilon_{eq}^S, \tag{8}$$

where for spin S, ε_{init}^S and ε_{eq}^S are the initial and equilibrium biases, t^S is the duration, and ${T_1}^S$ is the T_1 value. The following values were used:

- For C1: $\varepsilon_{init}^{\mathrm{C1}} = \varepsilon_{eq}^{\mathrm{H}} = 3.98 \pm 0.01$, $\varepsilon_{eq}^{\mathrm{C1}} = 1.000 \pm 0.003$, $t^{\mathrm{C1}} = t_1 + t_2$, $T_1^{\mathrm{C1}} = 43\mathrm{s} \pm 4\mathrm{s}$
- For C2: $\varepsilon_{init}^{\mathrm{C2}} = \varepsilon_{eq}^{\mathrm{H}} \left(1 - e^{-\frac{t1}{T_1^{\mathrm{H}}}}\right)$, $\varepsilon_{eq}^{\mathrm{C2}} = 1.000 \pm 0.003$, $t^{\mathrm{C2}} = t_2$, $T_1^{\mathrm{C2}} = 20\mathrm{s} \pm 2\mathrm{s}$
- For H: $\varepsilon_{init}^{\mathrm{H}} = 0$, $\varepsilon_{eq}^{\mathrm{H}} = 3.98 \pm 0.01$, $t^{\mathrm{H}} = t_2$, $T_1^{\mathrm{H}} = 3.5 \pm 0.1\mathrm{s}$

Biases and corresponding information contents (IC) were obtained for a fine mesh of delay combinations, in which each delay was varied between 0 and $5{T_1}^{\mathrm{H}}$ at 2 ms intervals. The resultant two-dimensional IC surface for the simplest model (assuming perfect PT steps) is shown in Figure 4; Note the broad plateau of near-maximal IC encompassing a large range of t_1 values. The maximal IC of 29.6 ± 0.7 was obtained with $t_1 = 9.604$s and $t_2 = 8.239$s, which yielded biases of 2.97 ± 0.08, 2.80 ± 0.08, and 3.60 ± 0.03, for C1, C2, and H, respectively. The errors were derived from Eq. 8 by the standard error formula

$$\Delta\varepsilon = \sqrt{\sum_i \left(\frac{\partial\varepsilon}{\partial x_i}\Delta x_i\right)^2},$$

where x_i are the relevant variables in each case.

Imperfect PT was incorporated into the simulation model by adding our experimental PT efficiencies (see Appendix C.1), namely $f_1 = 92\% \pm 2\%$ (initial H to C2), $f_2 = 69\% \pm 1\%$ (C2 to C1), and $f_3 = 74\% \pm 1\%$ (final H to C2). In Eq. 8, $\varepsilon_{init}^{\mathrm{C1}}$ was multiplied by $f_1 \cdot f_2$, while $\varepsilon_{init}^{\mathrm{C2}}$ was multiplied by f_3 (The errors formulas were adjusted accordingly). Application of this "practical" simulation model to POTENT yielded an optimal IC with the delays $t_1 = 11.032$s and $t_2 = 12.096$s; Note that the second delay is similar to the optimal experimental delay ($t_2 = 12$s), while the first delay is close (within one ${T_1}^{\mathrm{H}}$) to the optimal experimental value ($t_1 = 8$s). The practical simulation yielded an optimal IC of 22.3 ± 0.4 and corresponding biases of 1.89 ± 0.06, 1.99 ± 0.06, and 3.85 ± 0.02 (for C1, C2, and H, respectively). The same IC was obtained by applying the simulation model with the optimal experimental delays.

The IC obtained in our experiment was slightly lower, 20.7 ± 0.1. We attribute this small discrepancy to the naivete of our simulation model; the spins were assumed to relax independently, while we observed significant cross-relaxation during the extensive delays.

One may wonder whether the information content achieved by the two simulations should be considered as I — the information content, or $\tilde{I}$ — a lower bound on the information content. Under the extended Markov assumption, explained above, the correlations are assumed to cease to exist. Therefore the spins are in a tensor-product state, and the resulting biases lead directly to the information content I. If however, one questions the justification of the extended Markov assumption, and claims that correlations do remain, then the resulting information content calculated directly from the biases is $\tilde{I}$, a lower bound on the actual information content (in which part of the information remains in the correlations).

[5]The "extended-Markov model" is extremely useful in order to extend far beyond the naive ideal model of theortical AC — an infinite relaxation-times ratio (see more details in the post-scriptum).

The practical simulation model was also applied to the POTENT experiments that maximize 2-spin ICs, to determine whether bypassing the entropy bound was still feasible. Of particular interest was the case where only the two computation spins (carbons) were considered; Such a bypass would be impressive, as both spins would be significantly enhanced. Perfect PT allows significant bypasses considering any pair of spins, in particular, the two carbons are cooled to 3.30 ± 0.06, 3.82 ± 0.02, with a 2-spin IC of 25.5 ± 0.4, about 40% beyond the equilibrium IC of the spin system. The practical simulations indicate that the entropy bound can only be bypassed in the two (easier) cases where we ignore one of the carbons. In these cases the optimal 2-spin information contents are $I(\mathrm{C1},\mathrm{H}) \sim 20$, and $I(\mathrm{C2},\mathrm{H}) \sim 19$. We observed both bypasses experimentally (see Appendix C.2).

7 Discussion

We performed experimental effective-cooling of a spin system, implementing an essential step of algorithmic cooling,"heat-bath cooling". Furthermore, this is the first experiment designed specifically to bypass Shannon's bound on entropy manipulations, and (as far as we know) also the first one to actually bypass it. We quantitatively measured the extent to which the entropy bound was bypassed, showing a reduction of the total entropy by about 16%. Our dual-selective-reset experiment combined several polarization transfer steps, each between two selected spins, with relatively rapid thermal relaxation of the "reset spin"; The polarization of *two* ^{13}C nuclei was enhanced using *one* hydrogen in TCE.

Conventional PT techniques such as INEPT (and such as CP and NOE which are discussed in Appendix A) are all aimed at enhancing the signal of less polarized spins by transfer of polarization from nearby more polarized spins. The latter spin (usually 1H) typically relaxes faster than the target spin (commonly 13C or 15N). Hence, even for 2-spin systems (e.g., labeled chloroform) a combination of PT and a WAIT step is sufficient to obtain both a significantly cooler carbon and a repolarized proton, thus implementing a *single selective reset* — the most basic heat-bath cooling.

Let us consider the measurement results after a single PT and WAIT step. For example, an application of PT and WAIT to labeled chloroform could already bypass the entropy bound[6]. While we ourselves did not attempt such an "elementary bypass" using ^{13}C-chloroform, we implemented a quite similar elementary bypass on TCE by applying a single PT from the hydrogen to the nearby carbon, followed by WAIT. We bypassed Shannon's entropy bound, first by simulation (see section 6), and then also demonstrated it experimentally (see Appendix C). It is important to note that we are not aware of any previously reported elementary bypass (not even theoretically). Yet, interestingly, the somewhat simpler task of bypassing Sørensen's unitarity bound was discussed theoretically [24].

Going beyond such elementary bypasses achieved by a single-selective reset, the heat-bath cooling implemented here, and even more so — AC, achieve less trivial bypasses, since in these algorithms *pre-selected spins* are cooled via *multiple* selective reset. Thus, a main value of our work is in the *intentional* harnessing of selective reset steps for bypassing Shannon's entropy bound as a new effective-cooling method in NMR spectroscopy.

One question that might be asked is whether Shannon's entropy bound has ever been *unintentionally* bypassed in previous work. Various PT methods are often combined with rapid relaxation (e.g., by adding a paramagnetic salt) for improving standard signal averaging, by reducing the duration of each repetition [21, 25]. At certain moments, this combination may increase the polarization of the target spin, while also partially restoring the equilibrium polarization of the source spin, *possibly* bypassing the entropy bound. However, in standard signal averaging the measurement is not performed after an intermediate WAIT step, but rather after each PT step. Thus, whether anyone unintentionally bypassed the entropy bound in any past work will probably remain unknown.

One might ask to what extent the results of the current experiment point towards realistic near-future experiments in which more significant cooling might be achieved[7]. Our demonstration complements the previously implemented step of AC, polarization compression [26], and *together* with the polarization compression, our result demonstrates that AC could potentially become a practicable means for increasing spin polarization in NMR. Could this potential be fulfilled in the near future?

On the one hand, there are still several significant obstacles that need to be dealt with along the way, before significant cooling factors can be achieved using molecules that are relevant to some interesting applications in NMR spectroscopy. First, sufficiently large T_1 ratios are required in order to achieve more significant cooling. Secondly, as the gaps between the "ideal simulation" and the "practical simulation" (see Eq. 7) show, improving the PT efficiencies is also important in order to achieve more significant cooling. Last, implementation of the compression step is not trivial when the chemical shift between the computing spins is small, as in the case of the TCE molecule. See some relevant remarks in the Post-Scriptum.

On the other hand, a cooling factor of 1.5 or 2, relative to the cooling that is achievable today, could already be significant for various potential applications, e.g., for clinical NMR spectroscopy [27], and such a factor is well within the reach of practicable AC; unlike TCE, relevant biomolecules, such as amino acids, do have large chemical shifts, even at the typically low fields of clinical NMR spectroscopy (typically up to 3T). Furthermore, as we explain in the Post-Scriptum, high-fidelity gates have recently become available for suitable biomolecules containing 2 to 7 labeled carbons and one or more hydrogens.

Finally, we would like to briefly repeat the two connections to NMRQC (and to quantum computing in general): While the algorithms presented in [10, 13, 14, 18] are classical, they are implemented by quantum gates using tools developed in NMRQC, and therefore, our results (taken together with the experimental polarization compression [26]) hint that spin cooling for the purposes of magnetic resonance spectroscopy might provide the first near-future application of quantum computation techniques. In the long run, AC could also become relevant for generating nearly pure-state spins for NMR quantum computing; Use of electron spins in the future could lead to a breakthrough in experimental AC, as it would provide the high polarization of the electron, in addition to its very short reset time compared to nuclear spins.

[6] This could also be true for larger symmetric systems which contain several equivalent protons each bound to a carbon.

[7] This paragraph is rewritten in the postscriptum, taking into account more recent results.

8 Acknowledgements and Post-Scriptum

We thank Raymond Laflamme for helpful discussions and especially for participating in the initial stages of designing the experiment. We thank Yael Balasz, Sylvie Bilodeau, Jean-Christian Boileau, Nicolas Boulant, Camille Negrevergne and Tan Pham Viet for useful discussions, suggestions and help in setting up the experiments. We are grateful to Ilana Frank Mor for numerous helpful comments on the manuscript. The work of GB and JMF is supported in part by the Natural Sciences and Engineering Research Council of Canada. The work of GB is also supported in part by the Canada Research Chair programme and the Canadian Institute for Advanced Research. The work of TM, YW and YE is supported by the Israeli Ministry of Defense. The work of YE, HG, TM and YW is supported by the Institute for Future Defense Research at the Technion. JAJ and LX thank the UK Engineering and Physical Sciences Research Council and Biotechnology and Biological Sciences Research Council for financial support.

When we first submitted this paper to Science and then to PNAS in 2005 (See also [20]), the path to near-future applications seemed quite long, as we explain above. In hindsight, these results heralded further experiments, where more significant cooling was achieved. Our demonstration complemented the previously implemented step of AC, polarization compression [26], and was recently repeated with biomelecules [28]. The combination of both steps, initially in solid-state NMR, and recently also in liquid-state NMR [29, 30], demonstrates that AC could potentially become a practicable means for increasing spin polarization in NMR. Could this potential be fulfilled in the near future?

We wrote, in the first public version of this current paper [20], that the irreversible step of AC is first done in this work (which was correct then). However, shortly after this manuscript was submitted, a closely related work was described in [31] in which multiple selective reset steps were also done (independently of our work); In that work, AC was performed by combining multiple selective-reset steps with a compression step on a specially designed spin-system containing one active reset spin and three computer spins, using *solid-state* NMR. Their system contained ^{13}C labelled malonic acid in a single crystal containing mostly unlabelled malonic acid molecules. In that work, however, Shannon's bound was not bypassed. Subsequent work by the same group [32], published three years after we originally submitted the first revision of this paper describing our success at bypassing Shannon's bound, also bypassed the entropy bound. This was done using the same solid-state system they had previously used.

As expected, we recently achieved trivial bypasses with ^{13}C-chloroform at 500 MHz, e.g. the total entropy of this 2-spin system was reduced by about 20%, following PT from the proton to the carbon and a long delay of $\sim 7T_1(H)$, which allowed the proton to regain most of its equilibrium polarization (while the carbon retained a cooling factor of about 2).

Recent experimental "optimal control" methods in NMR [33, 34] are useful for generating high-fidelity gates. This progress, and the recent experimental heat-bath cooling of labeled amino acids [28], combined with recent theoretical algorithmic cooling results [35, 36], suggest that cooling ^{13}C spins in bio-molecules by a factor of ≈ 2 relative to the cooling which is achievable today, is likely within a few years. In particular, the compression step was already successfully implemented using optimal control tools both in solid-state NMR [32] and in liquid-state NMR[8]. In recent years AC was found useful in various directions [37–50] and some of the tools and results presented here could be relevant elsewhere as well.

[8] Experimental multi-cycle AC of labeled TCE, using optimal control, succeeded in bypassing the entropy bound on a single carbon, cooling it beyond $\sqrt{18}$. See more details in Appendix D, among additional Post-Scriptum details

References

[1] R. R. Ernst, G. Bodenhausen, and A Wokaun. *Principles of Nuclear Magnetic Resonance in One and Two Dimensions*. Oxford University Press, UK, 1987.

[2] C. P. Slichter. *Principles of Magnetic Resonance*. 3rd ed. Springer, Berlin, 1990.

[3] O. W. Sørensen. "Polarization transfer experiments in high-resolution NMR spectroscopy". In: *Prog. Nucl. Mag. Res. Spec.* 21 (1989), pp. 503–569.

[4] J. H. Ardenkjær-Larsen, B. Fridlund, A. Gram, G. Hansson, L. Hansson, M. H. Lerche, R. Servin, M. Thaning, and K. Golman. "Increase in signal-to-noise ratio of $> 10,000$ times in liquid-state NMR". In: *Proc. Natl. Acad. Sci. USA* 100 (2003), pp. 10158–10163.

[5] C. G. Joo, K. N. Hu, J. A. Bryant, and R. G. Griffin. "In situ temperature jump high-frequency dynamic nuclear polarization experiments: Enhanced sensitivity in liquid-state NMR spectroscopy". In: *J. Am. Chem. Soc.* 128 (2006), pp. 9428–9432.

[6] M. Anwar, D. Blazina, H. Carteret, S. B. Duckett, T. Halstead, J. A. Jones, C. Kozak, and R. Taylor. "Preparing high purity initial states for nuclear magnetic resonance quantum computing". In: *Phys. Rev. Lett.* 93 (2004).

[7] R. J. Fitzgerald, K. L. Sauer, and W. Happer. "Cross-relaxation in laser-polarized liquid xenon". In: *Chem. Phys. Lett.* 284 (1998), pp. 87–92.

[8] A. S. Verhulst, O. Liivak, M. H. Sherwood, H.-M. Vieth, and I. L. Chuang. "Non-thermal nuclear magnetic resonance quantum computing using hyperpolarized xenon". In: *Applied Physics Letters* 79.15 (2001), pp. 2480–2482.

[9] A. M. Oros and N. J. Shah. "Hyperpolarized xenon in NMR and MRI". In: *Phys. Med. Biol.* 49 (2004), R105–R153.

[10] L. J. Schulman and U. V. Vazirani. "Scalable NMR Quantum Computation". In: *ACM Symposium on the Theory of Computing (STOC): Proceedings*. 1999, pp. 322–329.

[11] R. B. Ash. *Information Theory*. Dover, New York, 1990.

[12] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley, New York, 1991, pp. 88–90.

[13] P. O. Boykin, T. Mor, V. Roychowdhury, F. Vatan, and R. Vrijen. "Algorithmic Cooling and Scalable NMR Quantum Computers". In: *Proc. Natl. Acad. Sci. USA* 99.6 (2002), pp. 3388–3393.

[14] J. M. Fernandez, S. Lloyd, T. Mor, and V. Roychowdhury. "Algorithmic Cooling of Spins: A Practicable Method for Increasing Polarization". In: *Int. J. Quant. Inf.* 2.4 (2004), pp. 461–467.

[15] T. Mor, V. Roychowdhury, S. Lloyd, J. M. Fernandez, and Y. Weinstein. *US patent No. 6,873,154*. 2005.

[16] D. G. Cory, R. Laflamme, E. Knill, L. Viola, T. F. Havel, N. Boulant, G. Boutis, E. Fortunato, S. Lloyd, R. Martinez, C. Negrevergne, M. Pravia, Y. Sharf, G. Teklemariam, Y. S. Weinstein, and W. Zurek. "NMR based quantum information processing: achievements and prospects". In: *Fortschr. Phys.* 48 (2000), pp. 875–907.

[17] J. A. Jones. "NMR Quantum Computation". In: *Prog. Nucl. Mag. Res. Spec.* 38.4 (2001), pp. 325–360.

[18] L. J. Schulman, T. Mor, and Y. Weinstein. "Physical limits of heat-bath algorithmic cooling". In: *Phys. Rev. Lett.* 94 (2005), p. 120501.

[19] J. M. Fernandez. "De computatione quantica". PhD thesis. University of Montreal, Canada, 2003.

[20] G. Brassard, Y. Elias, J. M. Fernandez, H. Gilboa, J. A. Jones, T. Mor, Y. Weinstein, and L. Xiao. "Experimental Heat-Bath Cooling of Spins". In: (2005). eprint: `arxiv:quant-ph/0511156v1`.

[21] G. A. Morris and R. Freeman. "Enhancement of nuclear magnetic resonance signals by polarization transfer". In: *J. Am. Chem. Soc.* 101 (1979), pp. 760–762.

[22] J. M. Fernandez, T. Mor, and Y. Weinstein. "Paramagnetic Materials and Practical Algorithmic Cooling for NMR Quantum Computing". In: *Int. J. Quant. Inf.* 3 (2005), pp. 283–285.

[23] H. K. Cummins, C. Jones, A. Furze, N. F. Soffe, M. M., J. M. Peach, and A. J. Jones. "Approximate Quantum Cloning with Nuclear Magnetic Resonance". In: *Phys. Rev. Lett.* 88 (2002), p. 187901.

[24] L. Emsley and A. Pines. "Lectures on Pulsed NMR (2^{nd} Ed.)" In: *Nuclear Magnetic Double Resonance, Proceedings of the CXXIII School of Physics "Enrico Fermi"*. World Scientific, Amsterdam, 1993, p. 216.

[25] R. Freeman. *Spin Choreography*. Oxford University Press, 1998.

[26] D. E. Chang, L. M. Vandersypen, and M. Steffen. "NMR implementation of a building block for scalable quantum computation". In: *Chem. Phys. Lett.* 338 (2001), pp. 337–344.

[27] T. B. Rodrigues and S. Cerdán. "13C MRS: An outstanding tool for metabolic studies". In: *Concepts Magn. Reson. Part A* 27A (2005), p. 1.

[28] Y. Elias, H. Gilboa, T. Mor, and Y. Weinstein. "Heat-bath cooling of spins in two amino acids". In: *Chem. Phys. Lett.* 517 (2011), pp. 126–131.

[29] Y. Atia. "Algorithmic Cooling of Spins by Optimal Control". MA thesis. Comp. Sci. Dept., Technion - Israel Institute of Technology, 2013.

[30] Y. Atia, Y. Elias, T. Mor, and Y. Weinstein. "Experimental Algorithmic Cooling in Liquid State NMR". manuscript in preparation.

[31] J. Baugh, O. Moussa, C. A. Ryan, A. Nayak, and R. Laflamme. "Experimental implementation of heat-bath algorithmic cooling using solid-state nuclear magnetic resonance". In: *Nature* 438 (2005), pp. 470–473.

[32] C. A. Ryan, O. Moussa, J. Baugh, and R. Laflamme. "Spin based heat engine: demonstration of multiple rounds of algorithmic cooling". In: *Phys. Rev. Lett.* 100 (2008), p. 140501.

[33] N. Khaneja, T. Reiss, C. Kehlet, S.-H. T., and S. J. Glaser. "Optimal control of coupled spin dynamics: design of NMR pulse sequences by gradient ascent algorithms". In: *J. Mag. Reson.* 172 (2005), pp. 296–305.

[34] Z. Tošner, T. Vosegaard, C. Kehlet, N. Khaneja, S. J. Glaser, and N. C. Nielsen. "Optimal control in NMR spectroscopy: Numerical implementation in SIMPSON". In: *J. Mag. Reson.* 197 (2009), pp. 120–134.

[35] L. J. Schulman, T. Mor, and Y. Weinstein. "Physical limits of heat-bath algorithmic cooling". In: *SIAM J. Comp.* 36 (2007), pp. 1729–1747.

[36] Y. Elias, T. Mor, and Y. Weinstein. "Semioptimal practicable algorithmic cooling". In: *Phys. Rev. A* 83 (2011), p. 042340.

[37] F Rempp, M Michel, and G Mahler. "A cyclic cooling algorithm". In: *Phys. Rev. A* 76 (2007), p. 032325.

[38] M. J. Henrich, F. Rempp, and G Mahler. "Quantum thermodynamic Otto machines: A spin-system approach". In: *Eur. Phys. J. Spec. Top.* 151 (2007), pp. 157–165.

[39] H Weimer, M. J. Henrich, F Rempp, H Schröder, and G Mahler. "Local effective dynamics of quantum systems: A generalized approach to work and heat". In: *Europhys. Lett.* 83.3 (2008), p. 30008.

[40] H. J Briegel and S. Popescu. "Entanglement and intra-molecular cooling in biological systems? - A quantum thermodynamic perspective". In: (2008). eprint: `arXiv:quant-ph/0806.4552v1`.

[41] N. Linden, S. Popescu, and P Skrzypczyk. "How small can thermal machines be? The smallest possible refrigerator". In: *Phys. Rev. Lett.* 105 (2010), p. 130401.

[42] O. C. O. Dahlsten, R. Renner, E. Rieper, and V. Vedral. "Inadequacy of von Neumann entropy for characterizing extractable work". In: *New J. Phys.* 13 (2011), p. 053015.

[43] W. S. Bakr, P. M. Preiss, M. E. Tai, R. Ma, J. Simon, and M. Greiner. "Orbital excitation blockade and algorithmic cooling in quantum gases". In: *Nature* 480 (2011), pp. 500–503.

[44] S. Simmons, R. M. Brown, N. V. A. Helge Riemann, P. Becker, H.-J. Pohl, M. L. W. Thewalt, K. M. Itoh, and J. J. L. Morton. "Entanglement in a solid-state spin ensemble". In: *Nature* 470 (2011), pp. 69–72.

[45] B. Criger, O. Moussa, and R. Laflamme. "Study of Multiple Rounds of Error Correction in Solid State NMR QIP". In: (2011). eprint: `arxiv:quant-ph/1103.4396`.

[46] R. Renner. "Thermodynamics: The fridge gate". In: *Nature* 482 (2012), pp. 164–165.

[47] A. Blank. "Scheme for a spin-based quantum computer employing induction detection and imaging". In: *Quantum Inf. Processing* 12 (2013), pp. 2993–3006.

[48] J. Baugh, J. Chamilliard, N. M. Chandrashekar, M. Ditty, A. Hubbard, R. Laflamme, M. Laforest, D. Maslov, O. Moussa, C. Negrevergne, M. P. da Silva, S. Simmons, C. A. Ryan, D. G. Cory, J. S. Hodges, and C. Ramanathan. "Quantum information processing using nuclear and electron magnetic resonance: review and prospects". In: (2007). eprint: `arxiv:quant-ph/0710.1447`.

[49] S. Lloyd. "Quantum optics: Cool computation, hot bits". In: *Nat. Photonics* 8 (2014), pp. 90–91.

[50] J.-S. Xu, M.-H. Yung, X.-Y. Xu, S. Boixo, Z.-W. Zhou, C.-F. Li, A. Aspuru-Guzik, and G.-C. Guo. "Demon-like algorithmic quantum cooling and its realization with quantum optics". In: *Nat. Photonics* 8 (2014), pp. 113–118.

[51] A. W. Overhauser. "Polarization of nuclei in metals". In: *Phys. Rev.* 92 (1953), p. 411.

[52] J. Granwehr, J. T. Urban, A. H. Trabesinger, and A. Pines. "NMR detection using laser-polarized xenon as a dipolar sensor". In: *J. Mag. Reson.* 176 (2005), p. 125.

[53] G. Navon, Y. Q. Song, T. Rōōm, S. Appelt, R. E. Taylor, and A. Pines. "Enhancement of Solution NMR and MRI with Laser-Polarized Xenon". In: *Science* 271 (1996), p. 1848.

[54] M. Anwar, J. A. Jones, D. Blazina, S. B. Duckett, and H. Carteret. "Implementation of NMR quantum computation with parahydrogen-derived high-purity quantum states". In: *Phys. Rev. A* 70 (2004), p. 032324.

[55] A. E. Derome. *Modern NMR Techniques for Chemistry Research*. Pergamon press, Oxford, UK, 1987.

[56] T. Mor. "Encyclopedia of Algorithms". In: Ed. Ming-Yang Kao. accepted for publication. Springer, Germany, 2007. Chap. "Algorithmic Cooling".

[57] Y. Elias, J. M. Fernandez, T. Mor, and Y. Weinstein. "Algorithmic Cooling of Spins". In: *Isr. J. Chem.* 46 (2006), pp. 371–391.

[58] D. P. Burum and R. R. Ernst. "Net polarization transfer via a J-ordered state for signal enhancement of low-sensitivity". In: *J. Mag. Reson.* 39 (1980), p. 163.

[59] G. Brassard, Y. Elias, T. Mor, and Y. Weinstein. "AC Limitations and Prospects". manuscript in preparation.

[60] Y. Atia, Y. Elias, T. Mor, and Y. Weinstein. "Quantum Computing Gates via Optimal Control". Submitted to Int. J. Quant. Inf.

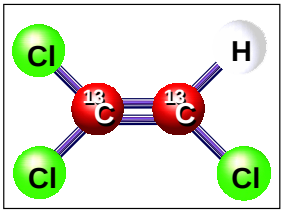

Figure 1: Trichloroethylene labeled with two ^{13}C. We denote the leftmost ^{13}C in this figure as C1 and the other, neighboring ^{1}H, as C2. In our experiments, the resonance frequencies were 125.773354, 125.772450 and 500.133245 MHz for C1, C2 and H, respectively. The scalar coupling constants were 201, 103 and 9 Hz between C2-H, C1-C2 and C1-H, respectively, while T_1 relaxation times were measured at 43 ± 4s and 20 ± 2s for C1 and C2, respectively, and 3.5 ± 0.1s for H.

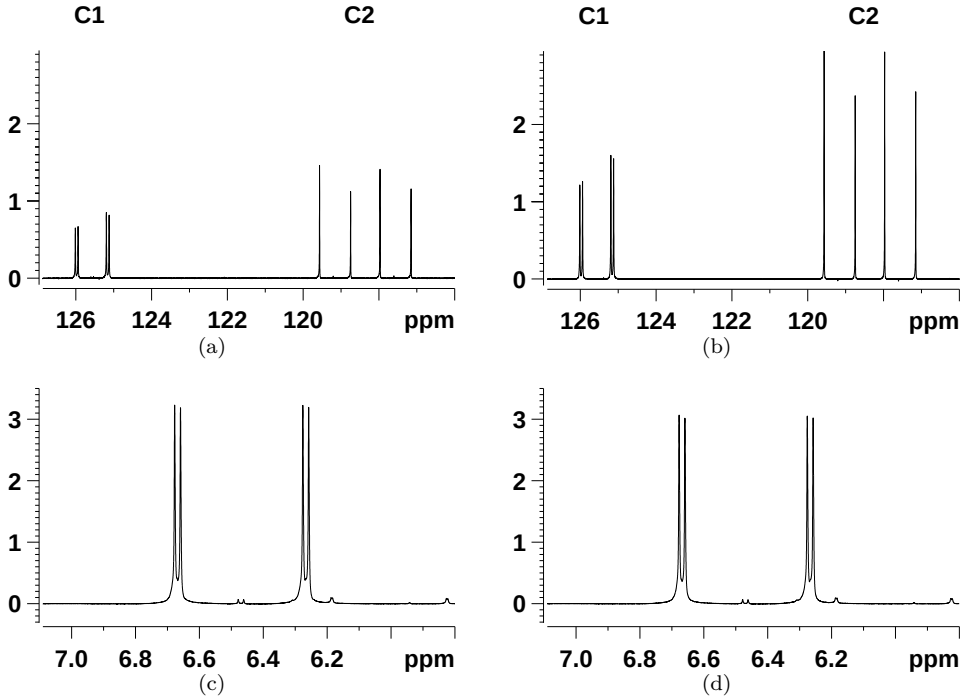

Figure 2: Spectra of TCE before and after the heat-bath cooling experiment. Vertical axes denote intensity in arbitrary units, scaled differently for each nucleus. Figs. (a) and (b) are the ^{13}C spectra before and after the experiment, respectively, with the left multiplet being C1 and the right one C2. Figs. (c) and (d) are the corresponding ^{1}H spectra before and after the experiment, respectively. The spectrum in Fig. (d) was obtained by running the AC experiment a second time with the exact same parameters as in Fig. (b), this time observing the ^{1}H instead of the ^{13}Cs by reversing the spectrometer channels.

A Effective-cooling approaches in NMR

When a spin-half particle is placed in a constant magnetic field and coupled to a thermal bath, the probabilities of finding it in the two spin states $\left|\pm\frac{1}{2}\right\rangle$ are given by the Boltzmann formula

$$P_{\pm\frac{1}{2}} = \frac{1 \pm \varepsilon}{2} = \frac{\exp(-E_{\pm\frac{1}{2}}/k_B T)}{Z},$$

where $E_{\pm\frac{1}{2}}$ are the energies of the two states, k_B the Boltzmann constant, T the bath temperature, Z a normalization factor (the partition function), and $\varepsilon \equiv P_{\frac{1}{2}} - P_{-\frac{1}{2}}$ is the population bias or polarization bias. The equilibrium population bias is therefore $\varepsilon = \tanh(\Delta E/2k_B T)$, where ΔE is the energy gap between the two states. At high temperatures this simplifies to $\varepsilon \xrightarrow{k_B T \gg \Delta E} \Delta E/2k_B T$. The equilibrium population bias in liquid-state NMR systems at room temperature is typically rather small — below 10^{-4}.

As a result of the small polarization bias, the obtained signal in dilute solutions at room temperature is often small if a single measurement (a single scan) is performed. Signal averaging over many scans is then required to attain sufficient signal-to-noise ratio. To reduce the experimental duration, a wide range of effective-cooling techniques were developed to increase population biases, without cooling the system. Some of these methods are briefly discussed below. The effects can be described in terms of the final bias achieved, or equivalently in terms of a spin temperature, defined by

$$T \triangleq \frac{\Delta E}{2k_B \tanh^{-1}\varepsilon} \xrightarrow{\varepsilon \ll 1} \frac{\Delta E}{2k_B \varepsilon}.$$

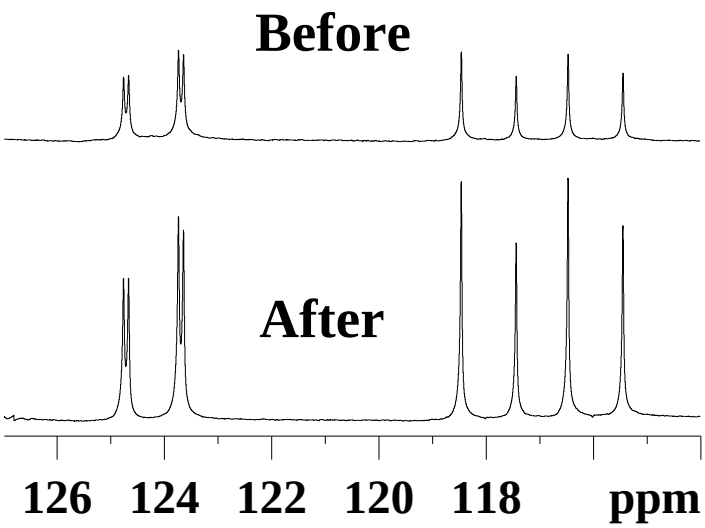

Figure 3: Carbon spectra before and after POTENT sequence from an experiment performed at Montréal in 2002.

Table 1: Initial and final polarizations, and final spin temperatures of each bit in TCE for the AC experiment shown in Figure 2.

spin	Initial bias (ε)	Sim bias	Prac-sim bias	Final bias	Final spin T (K)
C1 (far)	1.000 ± 0.003	2.87 ± 0.08	1.96 ± 0.06	1.74 ± 0.01	170 ± 1
C2 (adjacent)	1.000 ± 0.003	2.41 ± 0.09	1.90 ± 0.06	1.86 ± 0.01	159 ± 1
H	3.98 ± 0.01	3.85 ± 0.02	3.85 ± 0.02	3.77 ± 0.01	312 ± 1

The resonance frequency of each nucleus was used to compute its bias at thermal equilibrium at the room temperature of 296 ± 1 K in which the experiment was run. Precise values for the final polarization of each nucleus were obtained by comparing the integrals of the peaks, before and after the POTENT process; This method of calculation provides the final bias of each spin by tracing-out the other two spins (see [23]). The simulated values, in the third and fourth columns, were computed by taking into consideration empirical T_1 relaxation times. The results in the third column assume perfect polarization transfers, while the results in the fourth column use empirical transfer efficiencies.

Conventional polarization transfer. Elementary polarization transfer (PT) techniques such as INEPT are widely used in NMR studies [3], leading to modest polarization enhancements for the spins of interest. These techniques usually rely on transferring polarization from a hydrogen nucleus, with a relatively large bias, to a directly bonded carbon or nitrogen nucleus with a much smaller bias. These techniques are simple to implement, but cannot give large enhancements.

There are other techniques such as cross-polarization (CP), continuous CP, and NOE that can be used similarly to INEPT for transferring polarization from protons to carbons and other spins. In contrast with INEPT, where the PT is through bond, NOE is manifested through space, and is therefore simple to implement yet less efficient (limited to 75% of the proton polarization).

Dynamic Nuclear Polarization (DNP). Dynamic Nuclear Polarization [51], or DNP, refers to processes that transfer polarization from the spins of unpaired electrons to nearby nuclear spins, e.g. by saturating electronic transitions with microwave irradiation. DNP can yield polarizations close to 1 for solid samples at very low temperatures after prolonged irradiation. Much of the hyperpolarization is preserved when rapidly dissolving the sample [4], however dissolution DNP is slow and requires complex equipment, and therefore it cannot be used to refresh spins.

Electron Nuclear DOuble-Resonance (ENDOR). While not a spin-cooling approach *per se*, electron nuclear double-resonance (ENDOR) [2] takes advantage of the high polarization of electron spins to resolve NMR lines by detecting the resonances of nearby electron spins, e.g. in paramagnetic centers. The signal of a particular electron transition is first removed by saturating it with microwave radiation and then restored by irradiating at a radiofrequency around the hyper-fine coupling to exchange the population of electron and nuclear energy levels. The matching of electronic and nuclear transitions requires complex equipment and precise temperature control.

Hyperpolarized gases. Optical pumping approaches achieve extremely high spin polarizations in noble gases, most notably xenon [52]. A small portion of the hyperpolarization can be transferred to molecules of interest by co-condensing them with the noble gas and using simultaneous irradiation [7, 8, 53]. By this process, carbon and proton spins in chloroform were equally enhanced [8], and larger enhancements of about 20–40-fold were obtained at low temperature and field (200K and 1.4T) [7]. The transfer of polarization from xenon to ^{1}H or ^{13}C is non-selective and sometimes also non-uniform [7]. The process is complicated, requiring a laser array, an additional low magnetic field (0.1T), and an optical pumping chamber.

Parahydrogen induced polarization (PHIP). At low temperature (below about 85 K) in the presence of a suitable catalyst, hydrogen gas assumes its rotational ground state, in which the two hydrogen nuclei are in a spin

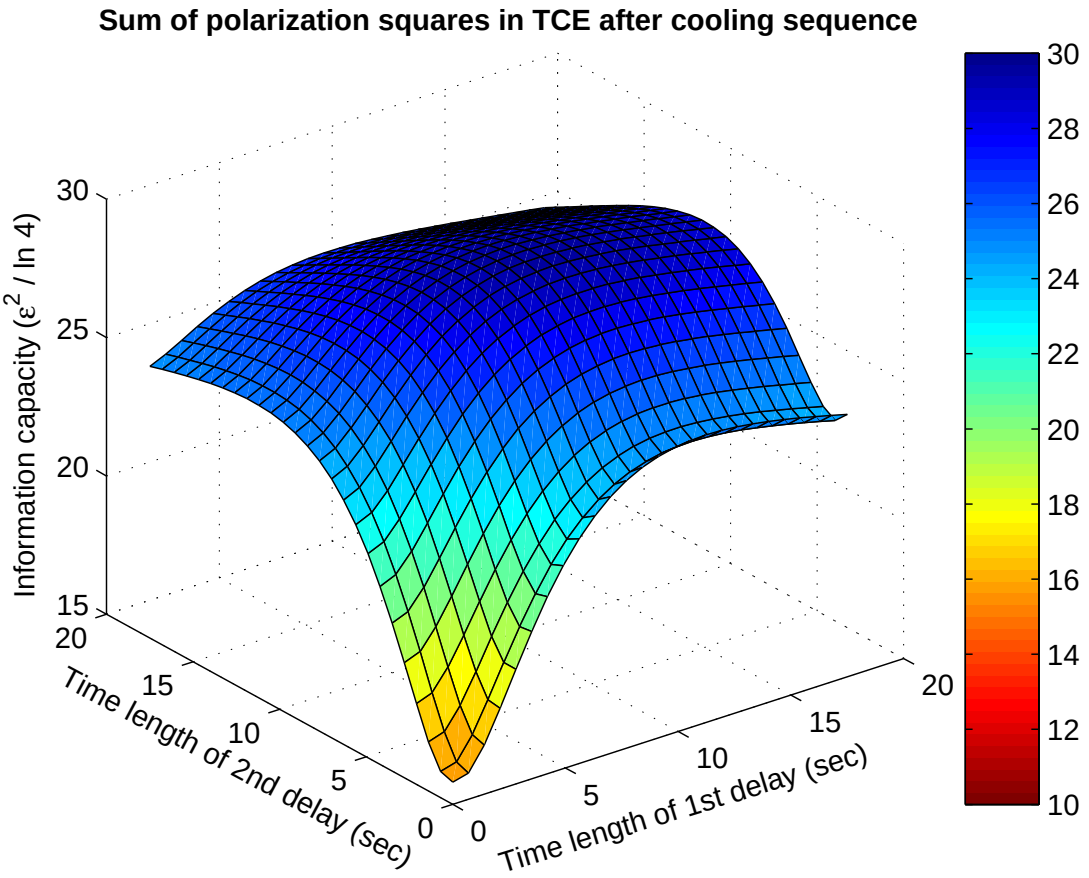

Figure 4: A simulation of information content (IC), labeled 'information capacity', as a function of the H repolarization delay times t_1 and t_2, with the IC represented on the z-axis. In this simulation we assume perfect polarization transfers. The maximum IC value was found numerically at $t_1 = 9.1$s–10.1s and $t_2 = 8.0$s–8.5s.

singlet state, known as *para*-hydrogen. This state is preserved upon warming to room temperature, and can be added across a double or triple bond [54]; the parahydrogen induced polarization (PHIP) is very high, sometimes approaching unity [54]. PHIP requires special equipment and cannot be generated in vivo.

Polarization compression and algorithmic cooling. In this appendix we focus so far on techniques that cool to the temperature of the high-bias spin. In the main text we describe going below the temperature of the high-bias spin using polarization compression and algorithmic cooling.

Beyond simple signal averaging. Last but not least, signal averaging may be enhanced by PT. First, when the source-spin for the PT is as polarized as the target spin (or even if it is less polarized than the target spin) by employing, as in algorithmic cooling, the idea of fast reset: Assuming the source and the target spins have the same polarization, but the source spin resets $\mathcal{R}$ times faster, PT is then simply used for obtaining a rapid recovery after each scan. Comparing this idea to algorithmic cooling was not done yet and is left for future research.

A similar case, in which the source spin is also useful due to its higher polarization, was demonstrated experimentally [21] and discussed in [55]. Due to the high relevance of this case to this paper, we discuss it in the "Discussion" section of the paper.

An important aspect of AC and its special case heat-bath cooling is the possibility of synergy with many other methods of signal enhancement: signal averaging, physical cooling, increasing the magnetic field, etc.

B Reversible compression and algorithmic cooling

B.1 From loss-less in-place data compression to polarization compression

Schulman and Vazirani (SV) [10] were the first to suggest that loss-less in-place data compression can lead to polarization compression. Consider a bit-string of length n, such that the probability distribution is known and far enough from the uniform distribution. One can use data compression to generate a shorter string, say of m bits, such that the entropy of each bit is much closer to one. As a simple example [56], consider a four-bit-string which is distributed as follows: $p_{0001} = p_{0010} = p_{0100} = p_{1000} = 1/4$, with p_i the probability of the string value i (the probability

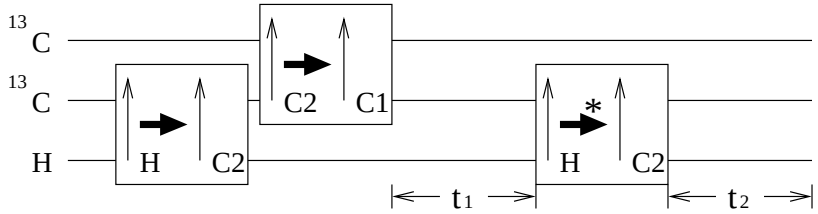

Figure 5: A block diagram of the complete experiment, with C1 represented at the top line. The arrow boxes denote polarization transfers in the direction of the arrow. The periods t_1 and t_2 are the variable delay times in which we wait for H to repolarize.

of any other value is zero). The bit-string can be compressed, via an algorithm, into a 2-bit string that holds the binary description of the location of "1" in the above four strings.

A similar loss-less, in-place, data compression process could generate an output of the same length n as the input, such that the entropy is compressed into the last two bits. For instance, logic gates that operate on the bits can perform the permutation, $\{0001 \rightarrow 0000; 0010 \rightarrow 0001; 0100 \rightarrow 0010; 1000 \rightarrow 0011\}$, while the other input strings (whose probability is anyhow zero) transform to output strings in which the two most significant bits are not zero; for instance $1100 \rightarrow 1010$. The entropy is now fully concentrated on the two least significant bits, thus the above process implements a simple case of data compression. The two most significant bits have zero entropy; if these two bits were spins, the process would have made them extremely cold.

In order to gain some intuition about the design of logic gates that perform such entropy manipulations when using nuclear spins, consider a closely related scenario (first considered by von Neumann): fair coin flips can be extracted, given a biased coin, by taking a pair of biased coin flips, with results a and b, and using the value of a conditioned on $a \neq b$. A simple calculation shows that $a = 0$ and $a = 1$ are now obtained with equal probabilities, and therefore the entropy of coin-flip a is increased in this case to 1; as we soon shall see, this means that its temperature is increased (to infinity). The opposite case, the probability distribution of a given that $a = b$, results in a highly determined coin flip; namely, a (conditioned) coin-flip with a higher bias or lower entropy. A gate that flips the value of b if (and only if) $a = 1$ is called a Controlled -NOT gate (CNOT). Consider the value of b following the CNOT, namely b_f; $b_f = 1$ implies that $a \neq b$ prior to the gate; the final entropy of a is then 1. On the other hand, $b_f = 0$ implies that $a = b$ prior to the gate; the final entropy of a in this case is lower than its initial value. The similar scenario of flipping two identical coins is more relevant to spin-entropy manipulations discussed below.

For quantum two-level systems (e.g., spin-half nuclei) there is a simple connection between temperature, entropy, and population probability. The process of increasing the polarization bias (reducing the entropy) without cooling the thermal-bath is known as "effective-cooling" (Appendix A). We can conclude that the two most significant bits in the first example got much colder (actually, to ZERO temperature) during data compression. The second example is directly relevant to nuclear spins, and we see that this process can cool some of them. The CNOT gate can be combined with other simple gates in a useful cooling subroutine.

SV [10] identified the importance of the low-entropy bits resulting from in-place loss-less data compression. Physical spin-half nuclei may be similarly cooled by data compression algorithms. Consider a linear molecule with 3 adjacent spin-half nuclei; At thermal equilibrium at room temperature and a constant magnetic field, the bits on each molecule are essentially uncorrelated. Furthermore, in the liquid state one can also neglect the interaction between molecules. It is convenient to write the probability distribution of a single spin at thermal equilibrium using the "density matrix" notation

$$\rho_\varepsilon = \begin{pmatrix} p_\uparrow & 0 \\ 0 & p_\downarrow \end{pmatrix} = \begin{pmatrix} (1+\varepsilon)/2 & 0 \\ 0 & (1-\varepsilon)/2 \end{pmatrix} , \tag{9}$$

since these two-level systems are of a quantum nature (namely, these are quantum bits — qubits), their states are not limited to a classical probability distribution over '0' and '1'. We consider now the classical case where ρ contains only diagonal elements that describe a conventional probability distribution. At thermal equilibrium, the state of $n = 2$ uncorrelated qubits with the same polarization bias is described by the density matrix $\rho_{\rm init}^{\{n=2\}} = \rho_\varepsilon \otimes \rho_\varepsilon$, where $\otimes$ means tensor product. The probability of the state '00', for instance, is $(1+\varepsilon)/2 \times (1+\varepsilon)/2 = (1+\varepsilon)^2/4$ (etc.) Similarly, the initial state of a 3-qubit system of this type, at thermal equilibrium, is $\rho_{\rm init}^{\{n=3\}} = \rho_\varepsilon \otimes \rho_\varepsilon \otimes \rho_\varepsilon$. This state represents a thermal probability distribution, such that the probability of the classical state '000' is $p_{000} = (1+\varepsilon_0)^3/2^3$, etc. In reality, spins generally have different initial biases, but as long as these differences are small (e.g., for homonuclear spins), we ignore them. Sufficient chemical shifts (resonance difference) for spins of each type are assumed, to allow selective addressing.

SV analyzed the cooling of such systems using various tools of data compression. Their technique and similar ones are now called "reversible polarization compression". Some ideas of SV were already explored a few years earlier by Sørensen [3], who analyzed effective-cooling of spins, with no connection to data compression. The entropy bound (due to Shannon) tells us that the total entropy of the system cannot be reduced by such processes. A tighter bound [3] applies when only unitary operations are allowed.

Example: The polarization bias of a single spin (bit) in an n-spin molecule at room temperature, assuming all polarization biases are initially ε and are much smaller than 1, cannot be reduced [14] by more than a multiplicative factor of $\sqrt{n}$: The total entropy of such a molecule, $H(n) = n(1-\varepsilon^2/\ln 4) + O(\varepsilon^4)$, is compressed so that $n-1$ spins have maximal entropy; the remaining spin satisfies at best $H(\text{single}) = 1 - (\sqrt{n}\varepsilon)^2/\ln 4 + O(\varepsilon^4)$, so that

$$\varepsilon_{final} \approx \sqrt{n}\varepsilon \tag{10}$$

(the approximation is valid as long as $\varepsilon_{final} \ll 1$). Spins actually cannot be cooled to the extent allowed by Shannon's

bound due to Sørensen's unitarity bound [3], since arbitrary entropy-preserving manipulations are not physically possible.

B.2 Algorithmic Cooling — a brief review

Boykin, Mor, Roychowdhury, Vatan and Vrijen suggested in 2002 a effective-cooling technique, which they named *Algorithmic Cooling (AC)* [13], or more specifically, heat-bath AC. Controlled interactions with a heat bath allow, theoretically, cooling much beyond entropy preserving processes. In order to pump entropy out of the system, AC employs designated computation spins together with rapidly relaxing reset spins, repeating the following steps several times: entropy compression, entropy shift onto reset spins using PT, and entropy removal from the system (reset spin repolarization).

The concept of AC led to practicable AC (PAC) [14] for cooling *small molecules*. PAC schemes use PT steps, reset steps, and a basic 3-spin compression step termed *3-bit-compression (3B-Comp)* [14] (based on the above-mentioned idea by von Neumann):

1. CNOT, spin B as a control and spin A as a target.
   Spin A is flipped if $B = 1$.
2. NOT on spin A.
3. CSWAP, spin A as a control, and spins B and C as targets.
   B and C are swapped if $A = 1$.

When the three spins have initial bias ε, this subroutine cools spin C: if $A = 1$ after the first step (and $A = 0$ after the second step), C is left unchanged (with its original bias ε); however, if $A = 0$ after the first step ($A = 1$ after the second step), spin B is cooled by a factor of about 2. The CSWAP places the new bias on C, which is therefore, on average, cooled by a factor of 3/2 (assuming biases much smaller than 1). We do not care about the biases of the other two spins, as they will undergo reset. One may consider the case in which those two remaining spins have reset spins as their neighbours. As a result, a single application of 3B-Comp, followed by reset of the two (potentially heated) spins by a simple PT from their neighboring reset spins cools the entire 5-spin system, once the reset spins go back to their initial equilibrium temperature.

The following details are provided here to clarify the theoretical differences between AC and polarization compression. We assume here that we are given a system of n spins, all with the same polarization bias, as in the case analyzed near the end of the previous section. In contrast to the case analyzed earlier, now one of the n spins is assumed to be a reset spin, and the remaining $n-1$ spins are computing spins. A PAC variant in which reset spins also participate in compression is called PAC2 in [14]. PAC2 is simple - all 3-bit compressions are always applied to spins with identical biases. On a 3-spin system (assuming the right-most is the reset spin), one gets the final biases $\{3/2, 1, 1\}$, in units of ε. On a 5-spin system (again, assuming the right-most is the reset spin), one gets the final biases $\{9/4, 3/2, 3/2, 1, 1\}$, where the $9/4 = (3/2)^2$ bias is obtained by applying 3B-Comp onto 3 spin already cooled to biases of 3/2. Generalizing this process to n spins one gets (ideally) biases of $\{(3/2)^{(n-1)/2}, \ldots, (3/2)^2, (3/2)^2, (3/2), (3/2), 1, 1\}$, in units of ε, for any odd number of spins (out of which, $n-1$ are computing spins and one is a reset spin). PAC2 proves an exponential advantage of AC over the best reversible cooling, as the latter can only improve the bias of the coldest spin by a factor of $\sqrt{n}$. In typical scenarios, reset spins have higher initial polarizations than computation spins; the bias amplification factor of $(3/2)^{(n-1)/2}$ is relative to the larger bias of the reset spin.

The cooling steps (reset and reversible polarization compression) can be repeated several times. Fernandez [19] considered two computation spins and one reset spin and analyzed *optimal cooling* of this system; By repeating the reset and compression exhaustively, the final biases of the three spins approach the limit of $\{2, 1, 1\}$ in units of ε, the equilibrium bias of the reset spin. Optimal AC [18, 57] leads to the exponential series: $\{...128, 64, 32, 16, 8, 4, 2, 1, 1\}$, so the coldest spin is cooled by a factor of 2^{n-2}.

B.3 Theoretical heat-bath cooling and algorithmic cooling of TCE

TCE can be ideally considered as a 3-spin system, with one reset spin and two computation spins. Its initial biases are $\{1, 1, 4\}$, with an information content (IC) of 18. Bypassing Shannon's entropy bound means increasing the IC of the 3-spin system above 18.

We show in the main text that an ideal POTENT leads to $\{4, 4, 4\}$, with an information content of 48, bypassing Shannon's entropy bound. A partial process, with just one selective reset, leads to the biases $\{1, 4, 4\}$ (or the biases $\{4, 1, 4\}$), with an information content of 33, which already bypasses Shannon's entropy bound. Note that this "single selective reset" contains only one PT and one reset (or a dual PT in case of transferring the polarization to the far carbon). Still, we did not find any paper reporting a bypass of the entropy bound via this process. POTENT, of course, is more powerful, as the repeated selective reset cools two computing spins, and can also lead further, to AC.

Some interesting (ideal) cooling processes with TCE are described below. By one selective reset step to the far carbon, plus one PT from the proton to the near carbon, the resulting biases become $\{4, 4, 0\}$ in the case of a PT (or $\{4, 4, 1\}$ in case of an ideal SWAP) with an information content of 32 (33). This process is important, since it focuses on cooling the two spins of interest, such that the system is cooled beyond the entropy bound. Our ideal simulation (see section 6) shows that it may be achieved given the relaxation times in our system. However, a practical simulation, accounting for realistic PT steps (see again, section 6), shows that our current experimental parameters do not support it, and more efficient PT steps are required. Our experimental efforts yielded cold carbons (see main text and Appendix C.2), but no bypass of the entropy bound in this interesting case; however, see the Post Scriptum for some very recent results.

Full AC can cool further. As explained above, an ideal PAC cools the far-carbon to $4 \times 3/2 = 6$, which bypasses the entropy bound already on a *single* spin, as bypassing $\sqrt{18}$ on that cooled spin is sufficient for that purpose. In the ideal case (biases of $\{4,4,4\}$), a perfect compression followed by resets of the two remaining spins would yield approximately (error in ε^3) $\{6,4,4\}$. The information content of this final state is 68, bypassing, by far, the entropy bound.

Interestingly, a perfect optimal compression applied onto biases of $\{4,4,4\}$ leads immediately to biases of $\{6,2,2\}$ and information content of 44, which is lower than the information content of the initial state (48) in that case; The residual information content is transfered into classical correlations (compression is unitary and uses classical gates, so the total entropy cannot change by that process). Our practical simulation shows that bypassing 4 on the colder carbon is not possible with our current system (see section 6).

A similar analysis, by the way, is relevant for understanding the differences between PT and SWAP: both are unitary and preserve entropy, but PT (e.g. INEPT) moves some of the information content into classical correlations, so that (for two spins) initial biases of $\{4,1\}$ change to final biases of $\{0,4\}$ by ideal PT, while they change to final biases of $\{1,4\}$ by ideal SWAP.

The optimal cooling of a single spin leads to biases of $\{8,0,0\}$ and an information content of 64. The optimal cooling of the entire 3-spin system leads to biases of $\{8,4,4\}$ and an information content of 96. Both require many reset steps and are thus rather impractical due to decoherence and imperfect PT and compression steps; however, even reaching biases around $\{5,3,3\}$ and an information content of 43 could be very interesting. See the Post Scriptum for some very recent results approaching such goals.

C The experiments

C.1 Experimental details

The initial polarization biases were calculated according to the high temperature approximation, $\varepsilon \approx \Delta E/2k_BT$ (see Appendix A), where $\Delta E = h\nu$, h and k_B are the Planck and Boltzmann constants, and T is the room temperature of 296K±1K. The resonance frequencies, ν, determined by averaging the resonance frequencies of the four lines of the multiplet of each spin, were $\nu^{\mathrm{H}} = 500.1332346 \cdot 10^6 \pm 0.9$ Hz, $\nu^{\mathrm{C1}} = 125.7735301 \cdot 10^6 \pm 0.7$ Hz, and $\nu^{\mathrm{C2}} = 125.7726234 \cdot 10^6 \pm 0.4$ Hz. The line-width of all lines was below 2 Hz (1.8 Hz, 1.4 Hz, and 0.8 Hz for H, C1, and C2, respectively). The resulting biases were $\varepsilon^{\mathrm{H}} = 4.05 \pm 0.01 \cdot 10^{-5}$ and $\varepsilon^{\mathrm{C1}} = \varepsilon^{\mathrm{C2}} = 1.020 \pm 0.003 \cdot 10^{-5}$; the error due to the high temperature approximation ($\varepsilon \approx \tanh(\varepsilon)$) was much smaller, in the ninth and tenth decimal place for the proton and carbons, respectively. The biases were normalized by setting the bias of C2, which was on resonance, to $\epsilon^{\mathrm{C2}} = 1.000$ with 3% error due to the uncertainty in room temperature. The error values in the initial biases were derived here by the standard error formula:

$$\Delta\epsilon = \sqrt{\left(\frac{\partial\epsilon}{\partial\nu}\Delta\nu\right)^2 + \left(\frac{\partial\epsilon}{\partial T}\Delta T\right)^2} = \epsilon\left[\sqrt{\left(\frac{\Delta\nu}{\nu}\right)^2 + \left(\frac{\Delta T}{T}\right)^2}\right].$$

A block diagram depicting the various stages of our experiment is shown in Figure 5. The first and third transfer sequences (both from H to C2) use different refocusing schemes; the latter was designed to retain the enhanced polarization of C1 by refocusing its evolutions (indicated in Figure 5 by $*$). The first two polarization transfers are implemented by two overlapping refocused-INEPT[9] sequences [58], as shown in Figure 6a and Figure 6b. After the second refocused-INEPT, a 90° pulse aligns the carbons along the z axis prior to the first WAIT (see Fig. 6c). The third refocused-INEPT is immediately followed by a selective inversion of C1, and prior to the second WAIT, the carbons are aligned along the z axis by a 90° pulse (see Fig. 6d). Finally, acquisition is performed measuring the free induction decay along the y axis[10] for either the carbons (shown in Fig. 6d), or the proton. During the second WAIT, the spins are aligned along the z axis, hence the need for a final 90° pulse, see Fig. 7.

Recall that a transfer of polarization can be bi-directional (achieved by implementing a SWAP gate), yet we chose to use a uni-directional PT, so that the resulting pulse sequence is more efficient. We observed that each INEPT-based PT sequences implements (up to some irrelevant phases) a dual-CNOT gate, namely CNOT(source,target)-CNOT(target,source), which is preferred here over a SWAP gate, that is equivalent to a triple-CNOT gate, namely CNOT(source,target)-CNOT(target,source)-CNOT(source,target). The PT pulse sequence is preferred over the SWAP pulse sequence since it is shorter and contains fewer basic pulses, say, about two-thirds.

The complete POTENT sequence, detailed in Figure 7, was run with various combinations of t_1 and t_2 delays. For statistics, several spectra were acquired under the same conditions. Reported values were obtained in five single-scan measurements (for each nucleus). In order to validate the sequence and estimate the transfer efficiencies, truncated versions of the complete pulse sequence were acquired, each version terminating at a different stage. Intermediate spectra obtained in this manner are shown in Figure 8.

We obtained efficiencies of $92\% \pm 2\%, 69\% \pm 1\%$, and $74\% \pm 1\%$, for the three PTs. These transfer efficiencies were obtained by comparing the peak integrals after truncated pulse sequences at the laboratory conditions of the full experiment (errors are due to uncertainty in the integrals).

C.2 Additional experimental results: cooling two spins

It would be of interest if the combined entropy of the carbons after truncated POTENT, that excludes the final WAIT step, would exceed the equilibrium entropy of the spin system. This would occur, for instance, if both biases increased

[9] The term refocused-INEPT coined by Burum and Ernst [58] is not connected to the term refocusing used elsewhere in this paper with respect to canceling out undesired evolutions.

[10] Equivalent to taking the trace of $\left(I^+ \equiv I_x + iI_y\right)\rho$, where ρ is the final density matrix.

three-fold, thereby cooling both carbons to one third of the room temperature. To ascertain whether this could be achieved, we performed experiments with the POTENT sequence with a very short second delay, about 2–3$T_2{}^*$, required for the elimination of coherences that cause undesired phases[11]. By setting the second delay, t_2, to 0.4s, the carbon re-heating is minimized. We reached biases of 2.03 ± 0.02 and 2.91 ± 0.04 for C1 and C2, respectively[12], yielding $I(\mathrm{C1}, \mathrm{C2}) = 12.6 \pm 0.2$. Although we did not succeed to bypass the entropy bound in this case, the two carbon spins were cooled quite significantly. These biases were obtained with $t_1 = 8$s, and the resulting temperatures were 145 ± 2K and 101 ± 1K for C1 and C2, respectively, as mentioned in the "Results" section of the main paper.

In experiments designed to achieve "elementary bypasses", namely, in "single-selective-reset" experiments, the combined information contents of the proton and the cooled carbon were sufficient to bypass the entropy bound. For C1 and H, we obtained biases of 1.61 and 3.95, with $t_1 = 2$s and $t_2 = 18$s, yielding $I(\mathrm{C1}, \mathrm{H}) = 18.17$. [In another experiment we obtained, for C1 and H, biases of 1.64 and 3.94, with $t_1 = 4$s and $t_2 = 18$s, yielding (again) $I(\mathrm{C1}, \mathrm{H}) = 18.17$.] For C2 and H, we obtained biases of 1.90 and 3.80, with $t_1 = 12$s and $t_2 = 12$s, yielding $I(\mathrm{C2}, \mathrm{H}) = 18.03$. [In another experiment we obtained, for C2 and H, biases of 1.79 and 3.85, with $t_1 = 12$s and $t_2 = 14$s, yielding (again) $I(\mathrm{C2}, \mathrm{H}) = 18.03$.] These results were obtained with no optimization, however they are already beyond the entropy bound, when considering the ± 0.1 error in information content obtained over five repetitions in the main experiment.

D Post-Scriptum details

Accounting for longitudinal relaxation. The limitation of algorithmic cooling due to finite ratios, $\mathcal{R}$, between the T_1 relaxation times of computation spins and the characteristic repolarization times τ of reset spins are discussed in Refs [36, 59]. The duration of each reset step is $T_{\mathrm{WAIT}} = d \cdot \tau$, where in liquid state NMR τ is T_1(reset), while in the solid state [31, 32] τ was the characteristic time for spin diffusion. In Ref [59] we analyze various cooling algorithms for several values of $\mathcal{R}$. Ideally, $T_1(\mathrm{comp}) \gg T_{\mathrm{run}} \gg T_{\mathrm{WAIT}} \gg \tau$, where $T_{\mathrm{run}} = T_{\mathrm{WAIT}} \cdot \mathrm{N}$ is the runtime of the algorithm and N is the number of reset steps. A partial analysis is given in Ref [36] for 2PAC. PAC2 [14] may ideally cool one spin in a 5-spin system that includes one reset spin by a factor of 2.25 after 9 reset steps [14, 57]. When $\mathcal{R} = 10000$ (similar to $\mathcal{R}$in Refs [31, 32]) and $d = 5$, the number of reset steps, $d \cdot \mathrm{N} = 5 \cdot 17 \ll \mathcal{R}$, and the cooling factor is essentially unchanged (2.23). However, when $\mathcal{R} \sim d \cdot \mathrm{N}$, the deviations become significant, e.g. the cooling factor reduces to 2.14 (1.81; 1.76) for $\mathcal{R} = 100(10; 5)$. In such cases, better cooling is obtained by choosing lower values of d; for example, when $\mathcal{R} = 10$ and $d = 3$, the cooling factor is improved to 1.85.

Heat-bath cooling of amino acid spin systems. In Ref [28], we applied heat-bath cooling to the ^{13}C labeled backbone of two amino acids, glutamate and glycine. Unlike TCE, the large chemical shift between the two labeled carbons allowed highly efficient 1 msselective pulses for excitation and inversion of either carbon. Sufficient T_1 ratios ($\mathcal{R} \sim 10$) were found between C1 and the alpha proton(s), However, the other T_1 ratio (involving C2) was quite low (around 2). Therefore, in order to reduce the total entropy of the spin system (beyond Shannon's bound) we applied a single selective reset such that only C1remained cool. In addition, we applied truncated POTENT sequences, where the final delay was omitted, to significantly cool both carbons (each by a factor of about 2.5). Similar results were obtained for glutamate under physiological conditions (temperature and pH).

Compression and AC of TCE. We attempted to perform the polarization compression step on TCE using non-selective pulses, as the small chemical shift did not permit efficient spin-selective pulses[13]. However, significant enhancement could not be obtained, due to accumulation of errors and extensive decoherence over the many (dozens) of pulses and delays. Highly efficient 3-bit compression was achieved in the solid state using quantum optimal control theory (the GRAPE algorithm) [31, 32]. We have recently [29, 30, 60] adopted this approach in the liquid state, and achieved short optimized 3-bit compression shaped pulses (around 15 ms) for TCE; the high efficiency of the compression ($\sim 90\%$) allowed us to cool one carbon of TCE by about 4.5-fold, beyond Shannon's bound, following several rounds of algorithmic cooling.

[11] $T_2{}^*$ denotes the apparent T_2, reflected in our experiments by the linewidth.

[12] For C1, the error was obtained by linear regression over 8 results obtained with $d2 = 1, 2, \ldots, 8$s ($R^2 = 0.999$), while for C2the error was obtained by logarithmic regression over 5 data points ($d2 = 4, 5, \ldots, 8$s, $R^2 = 0.995$).

[13] An earlier 3-bit compression experiment [26] employed spin-selective pulses and had very long coherence times (several seconds), yet the efficiency was still quite low (about 50%).

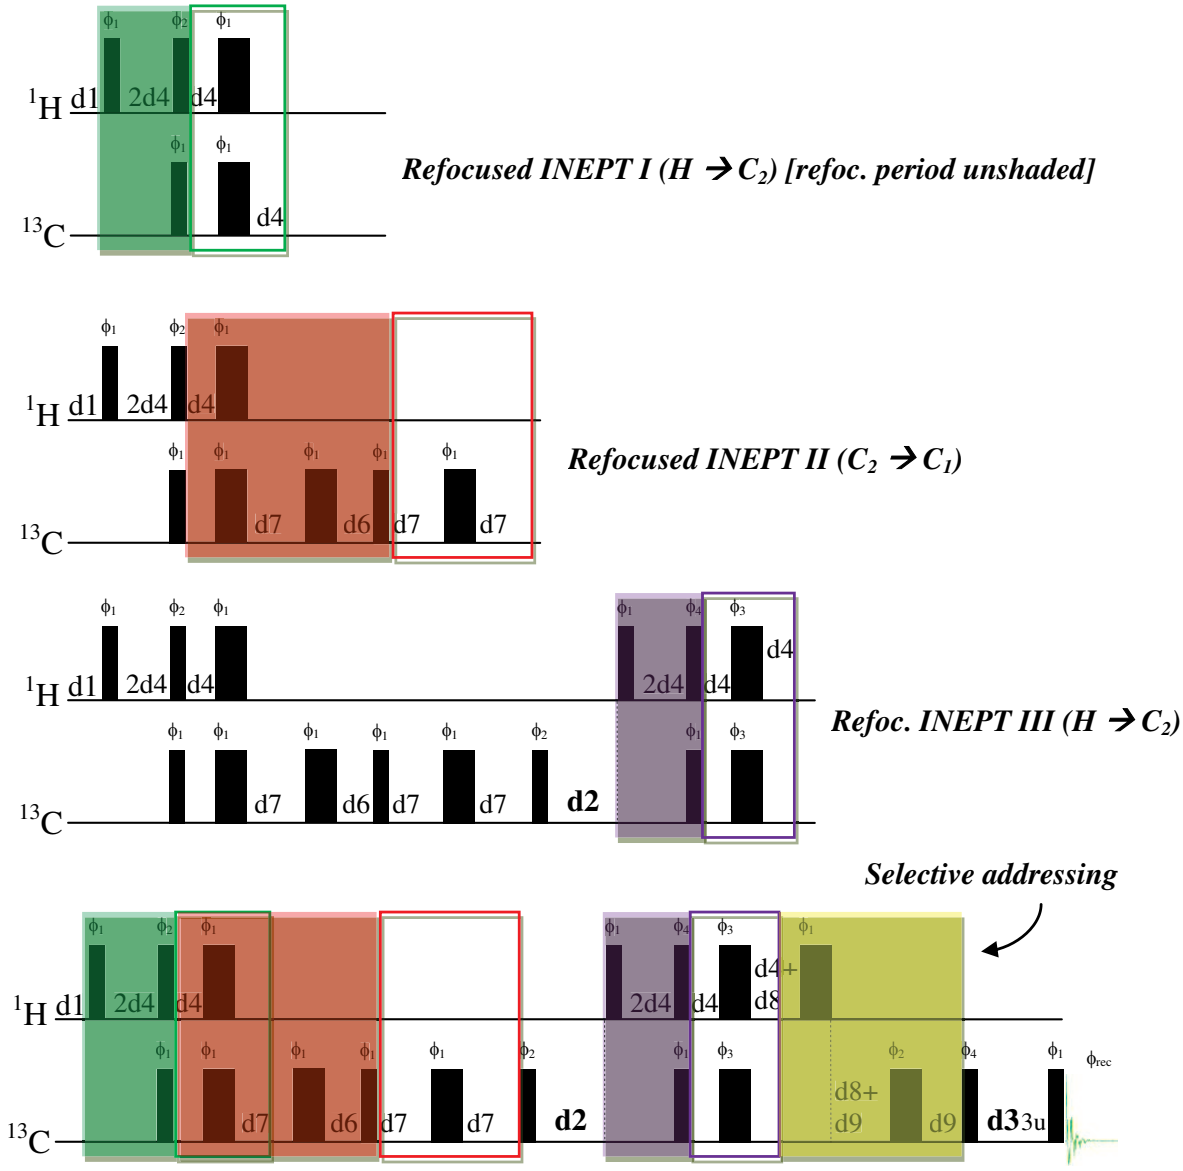

ϕ_1= x; ϕ_2= y; ϕ_3= -x; ϕ_4= -y ϕ_{rec}= x

d1=280s, **d2**=8s, **d3**=12s

d4=1/(4J_{CH}) = 1.24378 ms, d6=d7-d4 = 1.17987 ms, d7=1/(4J_{CC}) = 2.42366 ms

d8 = 1/(8δ_{CC}) = 137.614 µs, d9 = 2d7-d8-d4 = 3.46592 ms

Aq=2.99947s, 3u = 3µs

Figure 6: Annotation of the POTENT pulse sequence: (a), (b) and (c) highlight the first, second and third refocused-INEPTs; the first half is shaded, and the refocusing period is outlined by a rectangular border of the same color. In (d), the complete POTENT, cf. Fig. 7; the yellow block marks selective inversion of C1.

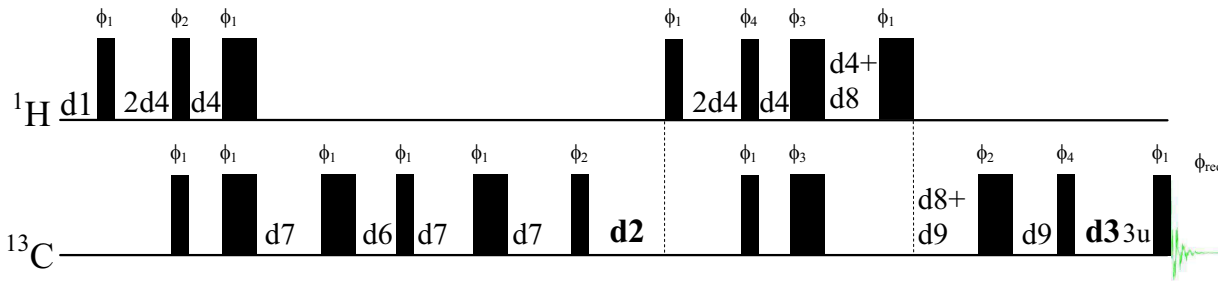

ϕ_1= x; ϕ_2= y; ϕ_3= -x; ϕ_4= -y ϕ_{rec}= x

Figure 7: The complete POTENT pulse sequence. Narrow bars and rectangles represent 90° and 180° rotations. See caption of Fig. 6 for values of the various delays. The Hrepolarization delays d2 and d3 are denoted in the text by t_1 and t_2.

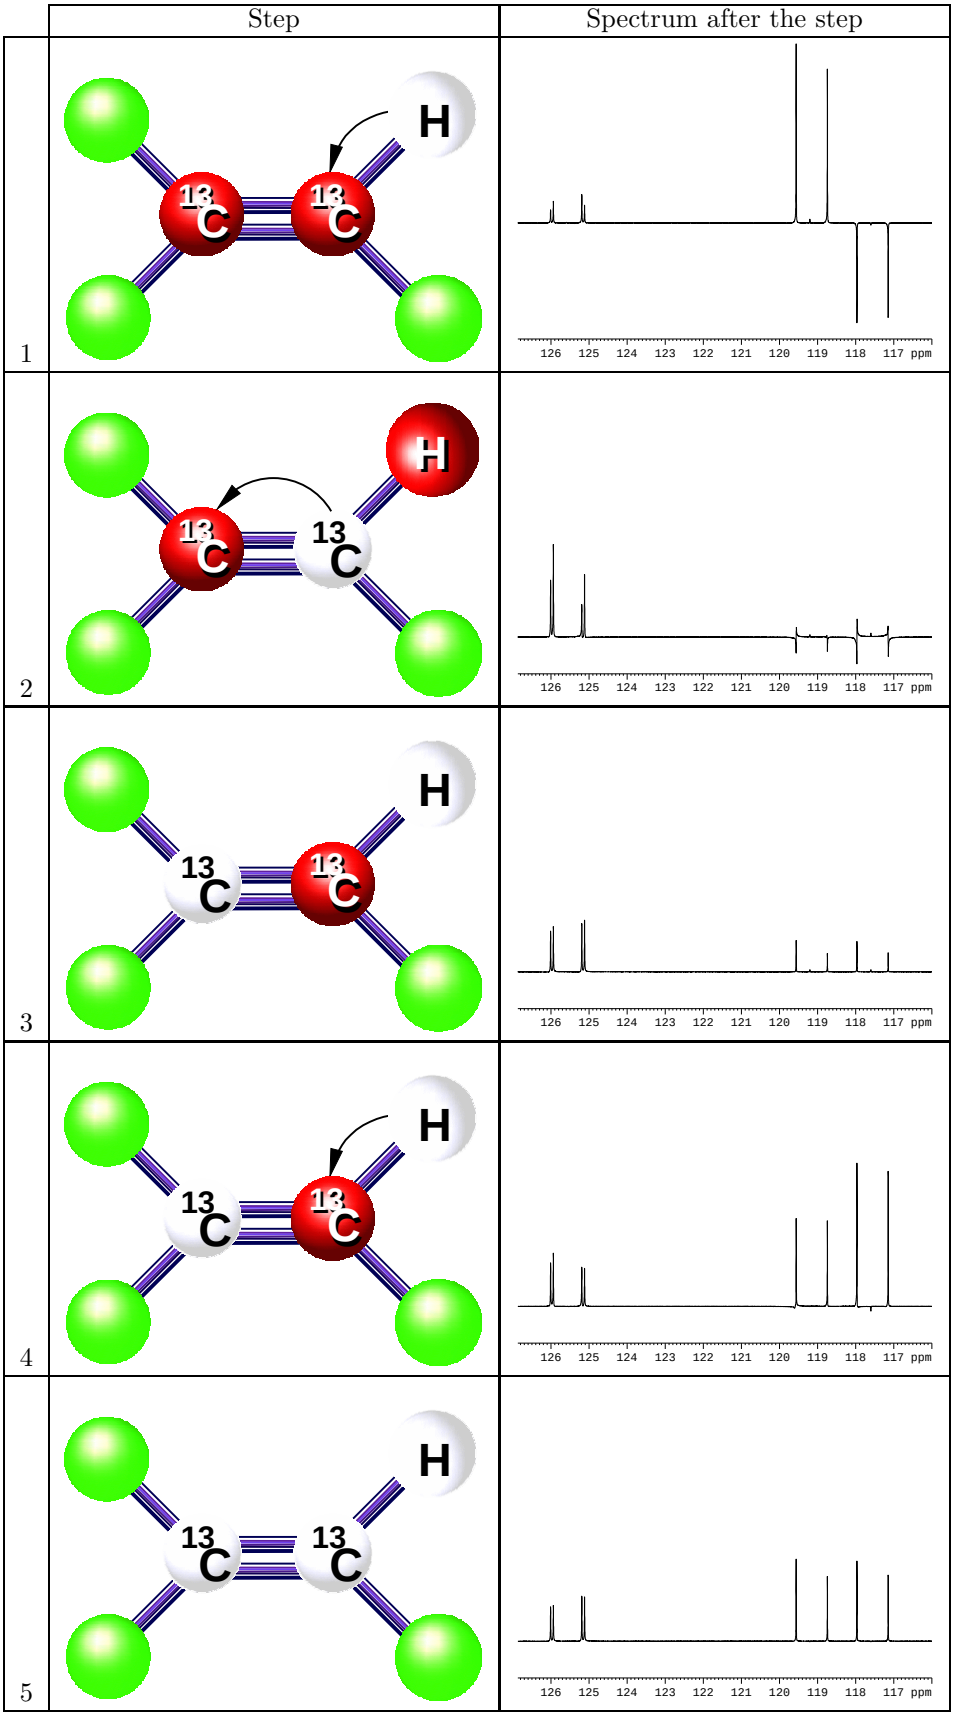

Figure 8: The steps of the cooling experiment and the resulting ^{13}C spectra after each step.